

TopicResponse: A Marriage of Topic Modelling and Rasch Modelling for Automatic Measurement in MOOCs

Jiazhen He^{1,3} · Benjamin I. P. Rubinstein¹ ·
James Bailey^{1,3} · Rui Zhang¹ ·
Sandra Milligan²

Received: date / Accepted: date

Abstract This paper explores the suitability of using automatically discovered topics from MOOC discussion forums for modelling students' academic abilities. The Rasch model from psychometrics is a popular generative probabilistic model that relates latent student skill, latent item difficulty, and observed student-item responses within a principled, unified framework. According to scholarly educational theory, discovered topics can be regarded as appropriate measurement items if (1) students' participation across the discovered topics is well fit by the Rasch model, and if (2) the topics are interpretable to subject-matter experts as being educationally meaningful. Such Rasch-scaled topics, with associated difficulty levels, could be of potential benefit to curriculum refinement, student assessment and personalised feedback. The technical challenge that remains, is to discover meaningful topics that simultaneously achieve good statistical fit with the Rasch model. To address this challenge, we combine the Rasch model with non-negative matrix factorisation based topic modelling, jointly fitting both models. We demonstrate the suitability of our approach with quantitative experiments on data from three Coursera MOOCs, and with qualitative survey results on topic interpretability on a Discrete Optimisation MOOC.

Keywords MOOCs · Topic Modelling · Matrix Factorisation · Psychometrics · Item Response Theory · Rasch Model

1 Introduction

Massive Open Online Courses (MOOCs) have attracted wide attention due to the promise of delivering education at scale. This new learning environment produces a

Jiazhen He
E-mail: jiazhenh@student.unimelb.edu.au

Benjamin I. P. Rubinstein, James Bailey, Rui Zhang and Sandra Milligan
E-mail: {brubinstein, baileyj, rui.zhang, s.milligan}@unimelb.edu.au

¹School of Computing & Information Systems, The University of Melbourne, Australia

²Melbourne Graduate School of Education, The University of Melbourne, Australia

³Data61/CSIRO, Australia

variety of data (*e.g.*, demographic data, student engagement, and forum activities), which offer new opportunities to understanding student learning. While quizzes and assignments have dominated summative assessment, the many sources of rich student engagement data generated in MOOC platforms present new views on student learning and avenues for formative feedback. This paper explores whether students' participation across automatically discovered MOOC forum topics is suitable for modelling academic ability.

Our work is inspired by the importance of forum discussions as an active learning activity, and recent research on quantitative measurement of student learning in the education community. In particular, i) MOOC discussion forums, as the main platform for student-instructor and student-student interactions, is of importance in gaining insights into student learning. ii) Recent research in education (Milligan, 2015) suggests that a distinctive and complex learning skill is required to promote learning in MOOCs. Educators are interested in whether and how the possession of this complex learning skill may be evidenced by latent complex patterns of engagement, instead of traditional assessment tools such as quizzes and assignments. iii) In order to validate such a hypothesis, measurement theory can be used (Rasch, 1993; Wright and Masters, 1982). A set of items is handcrafted from forum activities (*e.g.*, “*contributed a post attracting votes from others*” and “*made repeated thread visits in more than half the weeks*”), and calibrated (*e.g.*, deleted or changed) to fit a measurement model as evidence as to whether the set of items is appropriate for measuring the complex learning skill (Milligan, 2015). This process is human-intensive and time-consuming as reflected by Figure 1.

Driven by these observations, we investigate whether students' participation in automatically discovered forum topics can be used as an instrument to model students' ability. If students' participation across the discovered topics fit a measurement model (in this paper, we use the Rasch model) in terms of statistical effectiveness, and the topics are interpretable to subject-matter experts by way of qualitative effectiveness, then the discovered topics can be regarded as useful items for measurement. The resulting scaled topics, endowed with estimated difficulty levels, can assist in subsequent curriculum refinement, student assessment, and personalised feedback.

The technical challenge, then, is to automatically discover topics such that students' participation across them fit the Rasch model. He et al (2016) have adapted topic modelling of students' online forum postings, such that students' participation across these topics conforms to the Guttman scale. However, the Guttman scale is widely regarded as overly-idealised and impractical in the real world. In contrast the Rasch model, one of the simplest item response theory (IRT) models and the basis for many extensions, has been widely used in education and psychology. It is a generative probabilistic model that represents student responses as noisy observations of latent student abilities related to item difficulties. It can be viewed as a stochastic counterpart to the Guttman scale, permitting measurement error. If a person's ability level is higher than an item's difficulty, the person will answer the item correctly in the Guttman scale, while in the Rasch model there is a certain probability of incorrect response. While the Guttman scale only permits ordering of persons and items, Rasch models the locations on the scale and hence also meaningful differences (Scholten, 2011). The algorithm proposed for the Guttman scale (He et al, 2016) does not adapt readily for Rasch modelling. Instead we propose the TOPICRESPONSE algorithm, which simultaneously performs non-

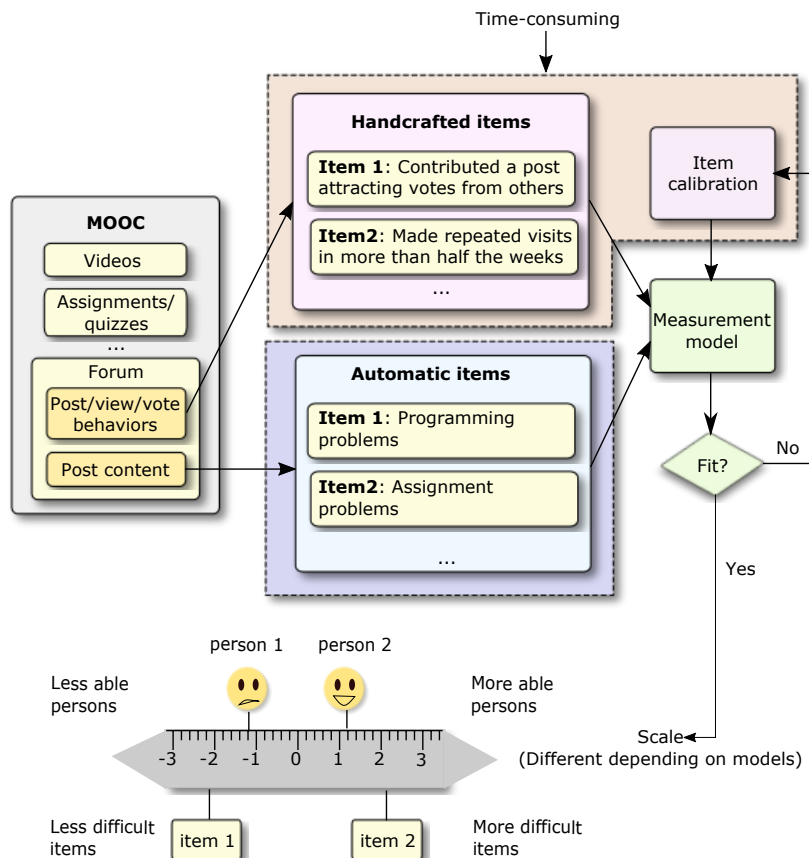


Fig. 1: Workflow for devising items manually versus automatically discovering topics as items for measurement. Traditionally, a set of items are handcrafted from MOOC forum behaviours, and then the students' dichotomous responses on the items are examined using the Rasch model. If the model fits well, then the students and items can be compared on an inferred scale (the ruler). Otherwise the items are refined (changed, added or deleted) manually until model fit. The process of handcrafting and calibration is time-consuming. Instead, we aim to automatically generate topics from discussion posts as items that fit the Rasch model by design.

negative matrix factorisation and Rasch model fitting. The main contributions of this paper include:

- The first study that combines topic modelling with Rasch modelling in psychometric testing: generating topics that measure students' academic abilities based on online forum postings;
- An algorithm TOPICRESPONSE fitting NMF and Rasch models simultaneously, for which we provide a proof of convergence; and
- Quantitative experiments on three Coursera MOOCs covering a broad swath of disciplines, establishing statistical effectiveness of our algorithm, and qualitative results on a Discrete Optimisation MOOC, supporting interpretability.

We review related work in Section 2. In Section 3, we present preliminaries and formalise our problem. Our algorithm is introduced in Section 4, and evaluated in Section 5. Section 6 concludes the paper.

2 Related Work

Many studies have focused on item response theory (IRT) or MOOC data analysis, but research on automatic discovery of items for measurement in MOOCs has received little attention. The main relevant work to this paper is (He et al, 2016), where NMF-based topic modelling is adapted and used for Guttman scaling (Guttman, 1950) in order to measure students’ latent abilities based on their MOOC forum posts. A major drawback of that work is that the Guttman scale is regarded to be the most restrictive IRT model and is overly idealised: it neither serves as the basis of more sophisticated (probabilistic) models, nor is it practical in the real world as a deterministic model. While the Guttman scale only models ordering of persons and items, the (probabilistic) Rasch model permits the interpretation of the differences between items and people (Scholten, 2011). The Rasch model is a generative model that models student responses as noisy observations of latent student abilities in relation to item difficulties. The algorithm for Guttman scaling (He et al, 2016) does not naturally extend to incorporating Rasch modelling.

2.1 Item Response Theory (IRT)

The field of IRT studies statistical models for measurement in education and psychology. Such models specify the probability of a person’s response on an item as a mathematical function of the person’s and item’s latent attributes. A principal goal of IRT is to create a scale on which persons and items can be placed and compared meaningfully. IRT has been used for computerised adaptive testing (CAT), which aims to accurately and efficiently assess individuals’ trait levels, and is used in the Scholastic Aptitude Test (SAT), Graduate Record Examination (GRE), while Chen et al (2005) proposed a personalised e-learning system based on IRT considering course material difficulty and learner ability.

As a statistical model, IRT has attracted attention in machine learning recently. Bergner et al (2012) applied model-based collaborative filtering to estimate the parameters for IRT models, considering IRT as a type of collaborative filtering task, where the user-item interactions are factorised into user and item parameters. Bachrach et al (2012) proposed a probabilistic graphical model that jointly models the difficulties of questions, the abilities of participants and the correct answers to questions in aptitude testing and crowdsourcing settings. While in MOOCs, Champaign et al (2014) investigated the correlations between resource use and students’ skill and relative skill improvement measured by IRT. Colvin et al (2014) analysed pre-post test questions using IRT, to compare the learning in MOOCs and a blended on-campus course. Past work has tended to focus on using already-devised items to measure student ability under IRT models, while we are interested in automatically discovering content-based items that are characteristic of measurement in MOOCs (Milligan, 2015).

2.2 MOOC Forums

MOOC forums have been of great interest recently, due to the availability of rich textual data and social behaviour. Various studies have been conducted such as sentiment analysis, community finding, question recommendation, answers & intervention prediction. Wen et al (2014) use sentiment analysis to monitor students' trending opinions towards the course and to correlate sentiment with dropouts over time using survival analysis. Yang et al (2015) predict students' confusion during learning activities as expressed in discussion forums, using discussion behaviour and clickstream data; they further explore the impact of confusion on student dropout. Ramesh et al (2015) predict sentiment in MOOC forums using hinge-loss Markov random fields. Gillani et al (2014) find communities using Bayesian Non-Negative Matrix Factorisation. Yang et al (2014) recommend questions of interest to students by designing a context-aware matrix factorisation model considering constraints on students and questions. MOOC forum data has also been leveraged in the task of predicting accepted answers to forum questions (Jenders et al, 2016) and predicting instructor intervention (Chaturvedi et al, 2014). Despite the variety of studies, little machine learning research has explored forum discussions for the purpose of measurement in MOOCs.

3 Preliminaries and Problem Formulation

We choose NMF as the basic approach to discover forum topics due to the interpretability of the topics produced, and the extensibility of its optimisation formulation. For the IRT model for measurement, we focus on the Rasch model for dichotomous data due to its popularity, and due to being the basis for many extensions in education and psychology. We next overview the Rasch model for dichotomous data and NMF, and then define our problem.

3.1 Rasch Model

The Rasch model (Wright and Masters, 1982; Bond and Fox, 2001) for dichotomous data (correct/incorrect, agree/disagree responses) specifies the probability of a person's positive response (correct, agree) on an item as a logistic function of the difference between the person's ability and item difficulty,

$$p_{ij} = P(X_{ij} = 1 | \beta_i, \theta_j) = \frac{1}{1 + \exp(-(\theta_j - \beta_i))} \quad (1)$$

where latent $\theta_j \in \mathbb{R}$ denotes person j 's ability, latent $\beta_i \in \mathbb{R}$ denotes item i 's difficulty, $X_{ij} \in \{0, 1\}$ denotes person j 's observed random response on item i , and p_{ij} is the probability of this response being positive. This probability is best illustrated with the Item Characteristic Curve (ICC) as depicted in Figure 2 and commonly used in the field of IRT. It can be seen that the higher a person's ability is, relative to the difficulty of an item, the higher the probability of a positive response on that item. When a person's ability is equal to an item's difficulty on the latent scale, positive responses are observed with 0.5 probability.

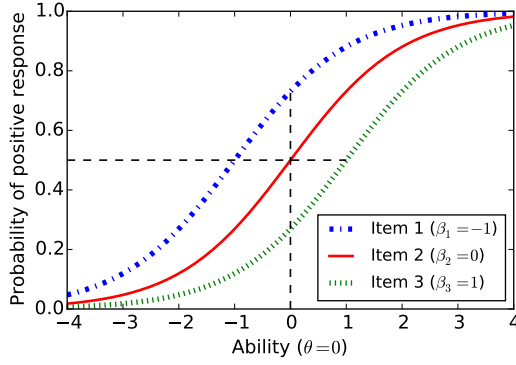


Fig. 2: The Item Characteristic Curves for three items (item 1—the easiest, 3—the most difficult). A person with ability $\theta = 0$ has 0.5 probability of responding positively on item 2 with difficulty $\beta = 0$, and higher (and lower) probability on the easiest item 1 (most difficult item 3, respectively).

The latent measurement scale is analogous to the ruler shown in Figure 1, where persons and items are placed together and can be compared meaningfully. The Rasch model provides a way to construct the ruler using persons' responses on items. Persons and items are located along the scale according to their abilities θ_j and difficulties β_i respectively.

The Rasch model can be viewed as a stochastic counterpart to the Guttman scale. For example, in Figure 1, person 1 and person 2 will have positive response on item 1 in a Guttman scale. While in a Rasch scale, there are certain probabilities that person 1 and person 2 will enjoy positive responses on item 1, with person 1's probability being higher. This error model leads to a higher level of measurement scale: the interval scale, where we can tell how much more able person 2 is compared to person 1. From the Guttman scale, by comparison, we can tell that person 2 is better than person 1 but not by how much.

Table 1 further illustrates our setup, with an example of items for measuring basic mathematical ability, alongside hypothetical students' responses. The initial estimates (see Equations 6,7 below) for item difficulties and person abilities are produced on a logit scale. For example, if person 1 responds to the items positively 20% of the time and negatively 80% of the time, then the person's initial ability estimate is approximately -1.39 by taking the natural logarithm of the odds ratio for positive response $\frac{0.2}{0.8}$.

3.1.1 Rasch Estimation

Given an observed response matrix $\mathbf{x}=[x_{ij}]$ (e.g., Table 1), a basic goal is to estimate the person and item parameters θ_j and β_i . The most common estimation methods are based on maximum-likelihood estimation, including: jointly maximum-likelihood (JML) estimation, conditional maximum-likelihood (CML) estimation and marginal maximum-likelihood (MML) estimation (Baker and Kim, 2004). In this paper, we focus on JML.

Table 1: An example of items for measuring basic mathematical ability, students' responses, initial item difficulty estimates and student ability estimates.

	Item 1 (Count)	Item 2 (+)	Item 3 (−)	Item 4 (×)	Item 5 (÷)	Proportion correct p_{θ_j}	Ability θ_j^0 $\log \left(\frac{p_{\theta_j}}{1-p_{\theta_j}} \right)$
Person 1	1	0	0	0	0	0.20	-1.39
Person 2	1	1	0	0	0	0.60	0.41
Person 3	0	1	1	0	0	0.60	0.41
Person 4	1	0	1	1	0	0.67	0.71
Person 5	1	1	1	0	1	0.80	1.39
Proportion correct p_{β_i}	0.80	0.33	0.33	0.20	0.20		
Difficulty β_i^0 $\log \left(\frac{1-p_{\beta_i}}{p_{\beta_i}} \right)$	-1.39	0.71	0.71	1.39	1.39		

Under the assumption that a sample of n persons is drawn independently at random from a population of persons possessing a latent skill attribute, and the assumption of local independence that a person's responses to different items are statistically independent, the probability of an observed data matrix $\mathbf{x} = [x_{ij}]$ with k items and n persons is the product of the probabilities of the individual responses, and can be given by the joint likelihood function

$$\begin{aligned} \mathcal{L}(\beta, \theta | \mathbf{x}) &= \prod_{i=1}^k \prod_{j=1}^n P(X_{ij} = 1 | \beta_i, \theta_j)^{x_{ij}} (1 - P(X_{ij} = 1 | \beta_i, \theta_j))^{(1-x_{ij})} \\ &= \prod_{i=1}^k \prod_{j=1}^n \frac{\exp(x_{ij}(\theta_j - \beta_i))}{1 + \exp(\theta_j - \beta_i)} . \end{aligned} \quad (2)$$

The log-likelihood function is then

$$\log \mathcal{L}(\beta, \theta | \mathbf{x}) = \sum_{i=1}^k \sum_{j=1}^n x_{ij}(\theta_j - \beta_i) - \sum_{i=1}^k \sum_{j=1}^n \log(1 + \exp(\theta_j - \beta_i)) . \quad (3)$$

The parameters of the Rasch model can be estimated by joint maximum likelihood—maximisation of this expression—using Newton-Raphson (Bertsekas, 1999), which yields the following iterative solution for β_i and θ_j ,

$$\beta_i^{t+1} = \beta_i^t - \frac{\sum_{j=1}^n (p_{ij} - x_{ij})}{-\sum_{j=1}^n p_{ij}(1 - p_{ij})} \quad \text{for } t \geq 0 , \quad (4)$$

$$\theta_j^{t+1} = \theta_j^t - \frac{\sum_{i=1}^k (x_{ij} - p_{ij})}{-\sum_{i=1}^k p_{ij}(1 - p_{ij})} \quad \text{for } t \geq 0 . \quad (5)$$

The convergence to a local optimum (with suitable step sizes) is guaranteed. The initial estimates of θ_j , θ_j^0 can be obtained by firstly calculating the proportion of items that a person j responded correctly p_{θ_j} , and then taking the natural

logarithm of the odds of person j 's correct response as shown in Table 1, which can be formalised as follows:

$$\theta_j^0 = \log \left(\frac{p_{\theta_j}}{1 - p_{\theta_j}} \right), \quad p_{\theta_j} = \frac{r_j}{k}, \quad r_j = \sum_{i=1}^k x_{ij}, \quad (6)$$

where r_j denotes the number of items that person j responded to positively. Similarly, the initial estimates of β_i, β_i^0 can be obtained by

$$\beta_i^0 = \log \left(\frac{1 - p_{\beta_i}}{p_{\beta_i}} \right), \quad p_{\beta_i} = \frac{s_i}{n}, \quad s_i = \sum_{j=1}^n x_{ij}, \quad (7)$$

where s_i denotes the number of persons who responded correctly on item i , and p_{β_i} denotes the proportion of persons who responded correctly on item i .

For those items receiving no correct responses ($s_i = 0$), or no incorrect responses ($s_i = n$), some implementations of the Rasch model will delete the item, while other models handle the situation as follows (Baker and Kim, 2004), where ϵ is a small number (*e.g.*, 1.0 is used in our experiments),

$$s_i = \begin{cases} \epsilon, & \text{if } s_i = 0 \\ n - \epsilon, & \text{if } s_i = n \end{cases}, \quad r_j = \begin{cases} \epsilon, & \text{if } s_i = 0 \\ k - \epsilon, & \text{if } s_i = k \end{cases}.$$

These pseudo counts are similar to frequentist Laplace corrections, or (weak) uniform Bayesian priors.

3.1.2 Evaluating Model Fit

A set of items is said to measure a latent attribute on an interval scale when there is a close fit between data and model. The model-data fit is typically examined using infit and outfit statistics—two types of mean square error statistics—conveying information about the error in the estimates for each individual item and person.

Outfit and infit test statistics are defined for each item and person to test the fit of items and persons under the Rasch model, by carefully summarising the Rasch residuals. The Rasch residuals are the differences between the observed responses and the expected responses according to the Rasch model. Formally, the expected response of person j on item i under the Rasch model $\mathbb{E}[x_{ij}]$ (abbreviated to E_{ij}) is $\mathbb{E}[X_{ij}] = p_{ij}$. The residual between the observation x_{ij} and the expected response E_{ij} is then $R_{ij} = x_{ij} - E_{ij}$. Standardised residuals are often used to assess the fit of a single person-item response

$$Z_{ij} = \frac{X_{ij} - E_{ij}}{\sqrt{\text{Var}(X_{ij} - E_{ij})}} = \frac{R_{ij}}{\sqrt{\text{Var}(X_{ij})}}, \quad (8)$$

where $\text{Var}(X_{ij}) = p_{ij}(1 - p_{ij})$ denotes the variance of X_{ij} (abbreviated to S_{ij}).

The outfit of item i summarises the squared standardised residuals, averaged over n persons,

$$\text{Outfit}_i = \frac{1}{n} \sum_{j=1}^n Z_{ij}^2 = \frac{1}{n} \sum_{j=1}^n \frac{R_{ij}^2}{S_{ij}}. \quad (9)$$

Typical treatments assume standardised residuals Z_{ij} approximately following a unit normal distribution. Their sum of squares therefore approximately follows a χ^2 distribution. Dividing this sum by its degrees of freedom yields a mean-square value, with an expectation of 1.0 and taking values in the range of 0 to infinity.

Outfit is sensitive to unexpected responses to items, *e.g.*, lucky guesses (*e.g.*, a person responds 111001) or careless sequences of mistakes (*e.g.*, a person responds 010100) (Linacre, 2002). Since outfit is sensitive to the very unexpected observations (outliers), infit was devised to be more sensitive to the overall pattern of responses (Linacre, 2006). Infit is an information-weighted form of outfit: it weights the observations by their statistical information (model variance) which is larger for targeted observations, and smaller for extreme observations (Bond and Fox, 2001). In this paper, we focus on infit. Formally, the infit of item i is given by

$$\text{Infit}_i = \frac{\sum_{j=1}^n S_{ij} Z_{ij}^2}{\sum_{j=1}^n S_{ij}} = \frac{\sum_{j=1}^n R_{ij}^2}{\sum_{j=1}^n S_{ij}} . \quad (10)$$

Both outfit and infit have the expected value of 1.0. Values larger than 1.0 indicate model underfitting, *i.e.*, data is less predictable than the model expects, while values less than 1.0 indicate overfitting, *i.e.*, observations are highly predictable (Wright et al, 1994). Conventionally, the acceptable range is usually taken to be $[0.7, 1.3]$ or $[0.8, 1.2]$ depending on application.

3.2 Non-Negative Matrix Factorisation (NMF)

Given a non-negative matrix $\mathbf{V} \in \mathbb{R}^{m \times n}$ and a positive integer k , NMF factorises \mathbf{V} into the product of a non-negative matrix $\mathbf{W} \in \mathbb{R}^{m \times k}$ and a non-negative matrix $\mathbf{H} \in \mathbb{R}^{k \times n}$ such that

$$\mathbf{V} \approx \mathbf{W}\mathbf{H}$$

\mathbf{V}
 $(m \times n)$

≈

\mathbf{W}
 $(m \times k)$

×

\mathbf{H}
 $(k \times n)$

.

A commonly-used measure for quantifying the quality of this approximation is the Frobenius norm between \mathbf{V} and $\mathbf{W}\mathbf{H}$. Thus, NMF involves solving

$$\underset{\mathbf{W}, \mathbf{H}}{\operatorname{argmin}} \|\mathbf{V} - \mathbf{W}\mathbf{H}\|_F^2 \quad \text{s.t.} \quad \mathbf{W} \geq \mathbf{0}, \quad \mathbf{H} \geq \mathbf{0} . \quad (11)$$

This objective function is convex in \mathbf{W} and \mathbf{H} separately, but not together. Therefore standard convex solvers are not expected to find a global optimum in general. The multiplicative update algorithm (Lee and Seung, 2001) is commonly used to find a local optimum, where \mathbf{W} and \mathbf{H} are updated by a multiplicative factor that depends on the quality of the approximation.

$$\begin{array}{c}
\begin{array}{ccccc}
& & \text{stud1} & \text{stud2} & \cdots & \text{stud}n \\
\text{solver} & \left(\begin{array}{cccc}
0.26 & 0.11 & \cdots & 0.52 \\
0.32 & 0.18 & \cdots & 0.06 \\
0.68 & 0.01 & \cdots & 0.83 \\
0.89 & 0.61 & \cdots & 0.44 \\
\vdots & \vdots & \ddots & \vdots \\
0.22 & 0.54 & \cdots & 0.98
\end{array} \right) \\
\text{optim} & \\
\text{code} & \\
\text{algorithm} & \\
\vdots & \\
\text{word}m &
\end{array} \\
\mathbf{V}
\end{array}
\quad
\begin{array}{c}
\begin{array}{ccccc}
& \text{topic1} & \text{topic2} & \cdots & \text{topic}k \\
\text{solver} & \left(\begin{array}{cccc}
0.22 & 0.01 & \cdots & 0.12 \\
0.38 & 0.15 & \cdots & 0.06 \\
0.18 & 0.05 & \cdots & 0.03 \\
0.09 & 0.21 & \cdots & 0.01 \\
\vdots & \vdots & \ddots & \vdots \\
0.02 & 0.04 & \cdots & 0.12
\end{array} \right) \\
\text{optim} & \\
\text{code} & \\
\text{algorithm} & \\
\vdots & \\
\text{word}m &
\end{array} \\
\mathbf{W}
\end{array}
\quad
\begin{array}{c}
\begin{array}{ccccc}
& \text{stud1} & \text{stud2} & \cdots & \text{stud}k \\
\text{topic1} & \left(\begin{array}{cccc}
0.83 & 0.17 & \cdots & 0.04 \\
0.21 & 0.75 & \cdots & 0.16 \\
\vdots & \vdots & \ddots & \vdots \\
0.09 & 0.64 & \cdots & 0.62
\end{array} \right) \\
\text{topic2} & \\
\vdots & \\
\text{topic}k &
\end{array} \\
\mathbf{H}
\end{array}
\end{array}$$

Fig. 3: Example matrices: word-student \mathbf{V} , word-topic \mathbf{W} , topic-student \mathbf{H} .

In the present MOOC setting, we focus on the students who contributed posts or comments in forums. For each student, we aggregate all posts or comments that they contributed. Each student is represented by a bag of words as shown in the example word-student matrix \mathbf{V} in Figure 3, where m represents the number of words, and n represents the number of students. Using NMF, a word-student matrix \mathbf{V} can be factorised into two non-negative matrices: word-topic matrix \mathbf{W} and topic-student matrix \mathbf{H} . For each student, the column vector of \mathbf{V} is approximated by a linear combination of the columns of \mathbf{W} , weighted by the components of \mathbf{H} . Therefore, each column vector of \mathbf{W} can be regarded as a topic, and the memberships of students in these topics are encoded by \mathbf{H} as shown in Figure 3.

3.3 Problem Statement

We seek to explore the feasibility of automatic discovery of forum discussion topics for measuring students' academic abilities in MOOCs, as quantified by the Rasch model. Our central tenet is that topics can be regarded as useful items for measuring a latent skill, if student responses to these topics are well fit by the Rasch model, and if the topics are interpretable to domain experts for educational relevance. Therefore, we need to discover topics from students' posts and comments in MOOC forums, in such a way that students' participation across these topics fits the Rasch model. Student item response records whether a student posts on the corresponding topic or not. After discovery, topics must then be further assessed for interpretability to domain experts. Our goal is decision support.

In particular, under the NMF framework, a word-student matrix \mathbf{V} can be factorised into two non-negative matrices: word-topic matrix \mathbf{W} and topic-student matrix \mathbf{H} . Our application requires that the topic-student matrix \mathbf{H} be **a)** binary ensuring the response of a student to a topic is dichotomous; **b)** useful for

Table 2: Glossary of symbols

Symbol	Description
m	the number of words
n	the number of students
k	the number of topics
$\mathbf{V} = (v_{ij})_{m \times n}$	word-student matrix
$\mathbf{W} = (w_{ij})_{m \times k}$	word-topic matrix
$\mathbf{H} = (h_{ij})_{k \times n}$	topic-student matrix
$\mathbf{H}_{ideal} = ((h_{ideal})_{ij})_{1 \times n}$	matrix for students with ideal number of distinct topics posted
$\mathbf{1}_r$	all-ones matrix with size $1 \times n$
g_j	student j 's grade
$\beta = (\beta_i)_k$	item difficulty vector
$\theta = (\theta_j)_n$	student ability vector
X_{ij}	binary response (0 or 1) of person j to item i
x_{ij}	observed response of person j to item i
p_{ij}	the probability of positive response of person j to item i
S_{ij}	variance of X_{ij}
Z_{ij}	standardised residual
$\lambda_0, \lambda_1, \lambda_2, \lambda_3$	regularisation coefficients

measuring students' academic abilities; and **c**) well-fit by the Rasch model. NMF provides an elegant framework for incorporating these constraints via adding novel regularisation, as detailed in the next section. A glossary of the symbols most used in this paper is given in Table 2.

4 The TOPICRESPONSE Algorithm: Joint NMF-Rasch Estimation

To favour topics that fit the Rasch model, we jointly optimise both NMF and Rasch models, which yields the objective function

$$g(\mathbf{W}, \mathbf{H}, \theta, \beta) = \|\mathbf{V} - \mathbf{WH}\|_F^2 - \lambda_0 f_R(\mathbf{H}, \theta, \beta) ,$$

$$f_R(\mathbf{H}, \theta, \beta) = \sum_{i=1}^k \sum_{j=1}^n h_{ij}(\theta_j - \beta_i) - \sum_{i=1}^k \sum_{j=1}^n \log(1 + \exp(\theta_j - \beta_i)) ,$$

where $f_R(\mathbf{H}, \theta, \beta)$ is the log-likelihood function maximised in Rasch estimation, and $\lambda_0 > 0$ is a user-specified parameter controlling the trade-off between the quality of factorisation and Rasch estimation.

Weak supervision of item responses. The fit between student topic responses \mathbf{H} and the Rasch model will provide statistical evidence of measuring skill attainment. However, it is difficult to conclude what the topics are measuring without domain knowledge. To favour the topics that can be used to measure students' academic abilities, we impose a constraint on \mathbf{H} based on some student grade, which provides an indicator of student's abilities (we discuss sources of auxiliary grade information below). In particular, we assume that there is the following relationship between

the ideal number of distinct topics that each student j contributes and their grade $g_j \in [0, 100]$,

$$(h_{ideal})_{1j} = \min \left\{ \left\lfloor \frac{g_j + width}{width} \right\rfloor, k - 1 \right\}, \quad width = \frac{100}{k - 1},$$

where \mathbf{H}_{ideal} is a $1 \times n$ matrix, denoting the ideal number of distinct topics posted by students. For example under $k = 10$ items, student j scoring $g_j = 45$ should post on a number of topics $(h_{ideal})_{1j} = 5$. The minimum and maximum number of different topics that a student j posted is 1 and $k - 1$ respectively. This is motivated by the initialisation of $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$ as illustrated in Section 3.1.1, where positive responses on 0 or k topics is undesirable.

This supervision constraint is markedly weaker than a similar constraint found in (He et al, 2016), as demonstrated in Figure 4. He et al (2016) leverage the student grade to exactly determine the item responses for the Guttman scale. The Guttman scale, as a deterministic model, requires that if a student can get a difficult item correct, they can also achieve correct responses on all easier items. This assumption is very restrictive, and rarely makes sense in practice. The Rasch model allows errors in the responses; and only constrains the number of distinct topics posted by a student, rather than the exact response pattern.

Most (MOOC) courses conduct multiple forms of assessment throughout the duration of teaching. For example, weekly quizzes, take-home assignments, mid-term tests, projects, presentations, etc. In the large-scale MOOC context, such evaluations may be peer-assessed. Students often enter courses with some cumulative grade-point average that may be (loosely) predictive of future performance. Any of these readily-available sources of student information could be reasonably used to seed \mathbf{H}_{ideal} . Even final course grades could be used, particularly when the ultimate application of TOPICRESPONSE is not measuring students, but refining curriculum.

In order to encourage satisfaction of the \mathbf{H}_{ideal} soft constraint on topic responses, we introduce a regularisation term on \mathbf{H} , namely $\|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2$.

	grade	8	25	46	67	89	98	78	35	55
\mathbf{H}_{ideal} for the Guttman scale		1	1	1	1	1	1	1	1	1
		0	1	1	1	1	1	1	1	1
		0	1	1	1	1	1	1	1	1
		0	0	1	1	1	1	1	1	1
		0	0	1	1	1	1	1	0	1
		0	0	0	1	1	1	1	0	1
		0	0	0	1	1	1	1	0	0
		0	0	0	0	1	1	1	0	0
		0	0	0	0	1	1	0	0	0
		0	0	0	0	0	1	0	0	0
		0	0	0	0	0	1	0	0	0
	grade	8	25	46	67	89	98	78	35	55
\mathbf{H}_{ideal} for the Rasch model		(1	3	5	7	9	9	8	4	5)

Fig. 4: An example of \mathbf{H}_{ideal} in the Guttman scale and the Rasch model.

Quantising & Regularising the Response Matrix. We introduce regularisation term $\|\mathbf{W}\|_F^2$, commonly used to prevent overfitting in NMF. To encourage binary solutions, we impose an additional regularisation term $\|\mathbf{H} \circ \mathbf{H} - \mathbf{H}\|_F^2$, where operator \circ denotes the Hadamard product. Binary matrix factorisation (BMF) is a variation of NMF, where the input matrix and the two factorised matrices are all binary. Our approach is inspired by those of Zhang et al (2007) and Zhang et al (2010). Our added term equals $\|\mathbf{H} \circ (\mathbf{H} - \mathbf{1})\|_F^2$, which is minimised by (only) binary \mathbf{H} .

TopicResponse Model. We have the following regularisations:

- $\|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2$ to encourage a grade-guided \mathbf{H} ;
- $\|\mathbf{W}\|_F^2$ to prevent overfitting; and
- $\|\mathbf{H} \circ \mathbf{H} - \mathbf{H}\|_F^2$ to encourage a binary item-response solution.

These terms together with joint NMF-Rasch estimation yield final objective

$$f(\mathbf{W}, \mathbf{H}, \boldsymbol{\theta}, \boldsymbol{\beta}) = \|\mathbf{V} - \mathbf{W}\mathbf{H}\|_F^2 - \lambda_0 f_R(\mathbf{H}, \boldsymbol{\theta}, \boldsymbol{\beta}) + \lambda_1 \|\mathbf{W}\|_F^2 + \lambda_2 \|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2 + \lambda_3 \|\mathbf{H} \circ \mathbf{H} - \mathbf{H}\|_F^2, \quad (12)$$

where $\lambda_1, \lambda_2, \lambda_3 > 0$ are user-specified regularisation parameters, with primal program

$$\underset{\mathbf{W}, \mathbf{H}, \boldsymbol{\theta}, \boldsymbol{\beta}}{\operatorname{argmin}} f(\mathbf{W}, \mathbf{H}, \boldsymbol{\theta}, \boldsymbol{\beta}) \quad \text{s.t.} \quad \mathbf{W} \geq \mathbf{0}, \mathbf{H} \geq \mathbf{0}. \quad (13)$$

TopicResponse Fitting Procedure. A local optimum of program (13) is achieved via iteration

$$w_{ij} \leftarrow w_{ij} \frac{(\mathbf{V}\mathbf{H}^T)_{ij}}{(\mathbf{W}\mathbf{H}\mathbf{H}^T + \lambda_0 \mathbf{W})_{ij}} \quad (14)$$

$$h_{ij} \leftarrow h_{ij} \frac{2(\mathbf{W}^T \mathbf{V})_{ij} + 8\lambda_2 h_{ij}^3 + 6\lambda_2 h_{ij}^2 + 2\lambda_1 (\mathbf{1}_r^T \mathbf{H}_{ideal})_{ij} + \lambda_3 (\boldsymbol{\theta} - \boldsymbol{\beta})_{ij}^+}{2(\mathbf{W}^T \mathbf{W}\mathbf{H})_{ij} + 12\lambda_2 h_{ij}^3 + 2\lambda_1 (\mathbf{1}_r \mathbf{H})_{ij} + 2\lambda_2 h_{ij} + \lambda_3 (\boldsymbol{\theta} - \boldsymbol{\beta})_{ij}^-} \quad (15)$$

$$\beta_i \leftarrow \beta_i - \frac{\sum_{j=1}^n (p_{ij} - h_{ij})}{-\sum_{j=1}^n p_{ij} (1 - p_{ij})} \quad (16)$$

$$\theta_j \leftarrow \theta_j - \frac{\sum_{i=1}^k (h_{ij} - p_{ij})}{-\sum_{i=1}^k p_{ij} (1 - p_{ij})} \quad (17)$$

where

$$\begin{aligned} (\boldsymbol{\theta} - \boldsymbol{\beta}) &= (\boldsymbol{\theta} - \boldsymbol{\beta})^+ - (\boldsymbol{\theta} - \boldsymbol{\beta})^- \\ (\boldsymbol{\theta} - \boldsymbol{\beta})_{ij}^+ &= \begin{cases} (\boldsymbol{\theta} - \boldsymbol{\beta})_{ij} & \text{if } (\boldsymbol{\theta} - \boldsymbol{\beta})_{ij} > 0 \\ 0 & \text{if otherwise} \end{cases} \\ (\boldsymbol{\theta} - \boldsymbol{\beta})_{ij}^- &= \begin{cases} -(\boldsymbol{\theta} - \boldsymbol{\beta})_{ij} & \text{if } (\boldsymbol{\theta} - \boldsymbol{\beta})_{ij} < 0 \\ 0 & \text{if otherwise} \end{cases} \end{aligned}$$

$(\boldsymbol{\theta} - \boldsymbol{\beta})^+$ and $(\boldsymbol{\theta} - \boldsymbol{\beta})^-$ denote the positive part and negative part of matrix $(\boldsymbol{\theta} - \boldsymbol{\beta})$ respectively. We next describe how these update rules are derived.

The update rules (16) and (17) can be obtained using Newton's method. The update rules (14) and (15) can be derived via the Karush-Kuhn-Tucker conditions necessary for local optimality. First we construct the unconstrained Lagrangian

$$\mathcal{L}(\mathbf{W}, \mathbf{H}, \boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\gamma}) = f(\mathbf{W}, \mathbf{H}, \boldsymbol{\theta}, \boldsymbol{\beta}) + \text{tr}(\boldsymbol{\alpha}\mathbf{W}) + \text{tr}(\boldsymbol{\gamma}\mathbf{H}) ,$$

where $\alpha_{ij}, \gamma_{ij} \leq 0$ are the Lagrangian dual variables for inequality constraints $w_{ij} \geq 0$ and $h_{ij} \geq 0$ respectively, and $\boldsymbol{\alpha} = [\alpha_{ij}]$, $\boldsymbol{\gamma} = [\gamma_{ij}]$ denote their corresponding matrices. The KKT condition of stationarity requires that the derivative of \mathcal{L} with respect to \mathbf{H} , vanishes at a local optimum $\mathbf{H}^*, \mathbf{W}^*, \boldsymbol{\alpha}^*, \boldsymbol{\gamma}^*$:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{W}} &= 2 \left(\mathbf{W}^* \mathbf{H}^* \mathbf{H}^{*T} - \mathbf{V} \mathbf{H}^{*T} + \lambda_0 \mathbf{W}^* \right) + \boldsymbol{\alpha}^* = \mathbf{0} , \\ \frac{\partial \mathcal{L}}{\partial \mathbf{H}} &= 2 \left(\mathbf{W}^{*T} \mathbf{W}^* \mathbf{H}^* - \mathbf{W}^{*T} \mathbf{V} + \lambda_1 \mathbf{1}_r^T \mathbf{1}_r \mathbf{H} + \lambda_2 \mathbf{H}^* - \lambda_1 \mathbf{1}_r^T \mathbf{H}_{ideal} \right) \\ &\quad + 4\lambda_2 \mathbf{H}^* \circ \mathbf{H}^* \circ \mathbf{H}^* - 6\lambda_2 \mathbf{H}^* \circ \mathbf{H}^* - \lambda_3 \left((\boldsymbol{\theta} - \boldsymbol{\beta})^+ - (\boldsymbol{\theta} - \boldsymbol{\beta})^- \right) + \boldsymbol{\gamma}^* \\ &= \mathbf{0} . \end{aligned}$$

Complementary slackness $\gamma_{ij}^* h_{ij}^* = 0$, implies:

$$\begin{aligned} 0 &= \left(\mathbf{V} \mathbf{H}^{*T} - \mathbf{W}^* \mathbf{H}^* \mathbf{H}^{*T} - \lambda_0 \mathbf{W}^* \right)_{ij} w_{ij}^* , \\ 0 &= \left(2\mathbf{W}^{*T} \mathbf{V} + 6\lambda_2 \mathbf{H}^* \circ \mathbf{H}^* + 2\lambda_1 \mathbf{1}_r^T \mathbf{H}_{ideal} - 2\mathbf{W}^{*T} \mathbf{W}^* \mathbf{H}^* - 2\lambda_1 \mathbf{1}_r^T \mathbf{1}_r \mathbf{H}^* \right. \\ &\quad \left. - 4\lambda_2 \mathbf{H}^* \circ \mathbf{H}^* \circ \mathbf{H}^* - 2\lambda_2 \mathbf{H}^* + \lambda_3 (\boldsymbol{\theta} - \boldsymbol{\beta})^+ - \lambda_3 (\boldsymbol{\theta} - \boldsymbol{\beta})^- \right. \\ &\quad \left. + 8\lambda_2 \mathbf{H}^* \circ \mathbf{H}^* \circ \mathbf{H}^* - 8\lambda_2 \mathbf{H}^* \circ \mathbf{H}^* \right)_{ij} h_{ij}^* . \end{aligned}$$

These two equations lead to the updating rules (14) and (15). Regarding the update rules (14), (15), (16) and (17) we have the following theorem:

Theorem 1 *The objective function $f(\mathbf{W}, \mathbf{H}, \boldsymbol{\theta}, \boldsymbol{\beta})$ of TOPICRESPONSE program (13) is non-increasing under update rules (14), (15), (16) and (17).*

This result guarantees that the update rules of \mathbf{W} , \mathbf{H} , $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$ eventually converge, and that the obtained solution will be a local optimum. The proof of Theorem 1 is given in the Appendix.

Algorithm 1 TOPICRESPONSE

Require:

\mathbf{V} , \mathbf{H}_{ideal} , λ_0 , λ_1 , λ_2 , λ_3, k ;

Ensure:

A topic-student matrix, \mathbf{H} , item difficulties $\boldsymbol{\beta}$, person abilities $\boldsymbol{\theta}$;

- 1: Initialise \mathbf{W} , \mathbf{H} using NMF;
 - 2: Normalise \mathbf{W} , \mathbf{H} following (Zhang et al, 2007, 2010);
 - 3: Initialise $\boldsymbol{\theta}$, $\boldsymbol{\beta}$ based on Eq. (6) and Eq. (7);
 - 4: **repeat**
 - 5: Update $\mathbf{W}, \mathbf{H}, \boldsymbol{\beta}, \boldsymbol{\theta}$ iteratively based on Eq. (14) to Eq. (17);
 - 6: **until** converged
 - 7: **return** \mathbf{H} ;
-

Algorithm. Our overall approach TOPICRESPONSE is described as Algorithm 1. \mathbf{W} and \mathbf{H} are initialised using plain NMF (Lee and Seung, 1999, 2001), then normalised (Zhang et al, 2007, 2010). $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$ are initialised based on Eq. (6) and Eq. (7), where x_{ij} is replaced by h_{ij} . At optimisation completion, estimates for topics, item difficulties and person abilities can be obtained together. Code for TOPICRESPONSE is available from the authors’ websites.

5 Experiments

We report on extensive experiments evaluating the effectiveness of TOPICRESPONSE on real MOOCs. In our experiments, we use the first offerings of three Coursera MOOCs from education, economics and computer science offered by The University of Melbourne: *Assessment and Teaching of 21st Century Skills* delivered in 2014, *Principles of Macroeconomics* delivered in 2013, and *Discrete Optimisation* delivered in 2013. We denote these three courses by EDU, ECON and OPT respectively.

5.1 Dataset Preparation

We focus on the students who contributed posts or comments in forums. For each student, we aggregate all the posts and comments that they contributed. After stemming and removing stop words, a word-student matrix with normalised tf-idf in $[0,1]$ is produced. The statistics of words and students before and after preprocessing, the dominated words, and the sparsity of word-student matrix (the percentage of non-zeros values) for three MOOCs are displayed in Table 3.

5.2 Baseline and Evaluation Metrics

We compare our algorithm TOPICRESPONSE with the baseline algorithm Grade-Guided NMF (GG-NMF), which minimises the following objection function

Table 3: Statistics of our three Coursera MOOC datasets.

MOOC	#Students	#Words	#Words after preprocessing	Dominated words	Word-student matrix sparsity
EDU	1,749	28,931	18,391	student, learn, skill, work, teacher, use, assess, teach, problem, collabor	0.59%
ECON	1,551	26,370	21,412	gdp, would, econom, think, product, good, one, economi, increas, invest	0.50%
OPT	1,092	19,284	16,128	use, solut, get, time, one, tri, python, work, optim, would	0.85%

$$f_G(\mathbf{W}, \mathbf{H}) = \|\mathbf{V} - \mathbf{WH}\|_F^2 + \lambda_1 \|\mathbf{W}\|_F^2 + \lambda_2 \|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2 + \lambda_3 \|\mathbf{H} \circ \mathbf{H} - \mathbf{H}\|_F^2$$

A local optimum can be obtained using the Karush-Kuhn-Tucker conditions. Like TOPICRESPONSE, GG-NMF regularises \mathbf{H} by considering the students' grades as an indicator of academic ability. The difference is that TOPICRESPONSE optimises the Rasch estimation and NMF simultaneously, while in GG-NMF, the students' topic responses \mathbf{H} are first obtained, and then are passed through the Rasch model. We evaluate the two algorithms in terms of the following metrics.

Quality of factorisation. We measure $\|\mathbf{V} - \mathbf{WH}\|_F^2$ so as to record how well the factorisation approximates the student-word matrix.

Measuring student academic ability. Quality of constraint on students' topic participation, based on grades: $\|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2$.

Negative log-likelihood. Log-likelihood measures fit of the Rasch model to the entire dataset. For convenience, we look at the negative log-likelihood, which should be minimised: smaller is better. This measure is our main focus for Rasch, as it is important to examine the model-level fit before looking at item-level fit.

Item infit. As illustrated in Section 3.1.2, item infit examines the fit of a particular item, with non-fitting items suitable for further refinement. We use the conventional acceptance range of [0.7, 1.3].

5.3 Hyperparameter Settings

Table 4 presents the parameter values used for our parameter sensitivity experiments, where the default values shown in boldface are used in experiments unless noted otherwise.

5.4 Main Results for GG-NMF and TopicResponse

In the first group of experiments, we examine the performance of GG-NMF (baseline) and TOPICRESPONSE in terms of negative log-likelihood, the quality of factorisation \mathbf{WH} in approximation \mathbf{V} given by $\|\mathbf{V} - \mathbf{WH}\|_F^2$, and the supervision

Table 4: Hyperparameter settings.

Param.	Values Explored (Default)
λ_0	[0.01, 0.1 , 0.2, 0.3, 0.4, 0.5]
λ_1	[10^{-3} , 10^{-2} , 10^{-1} , 10^0 , 10^1 , 10^2]
λ_2	[10^{-3} , 10^{-2} , 10^{-1} , 10^0 , 10^1 , 10^2]
λ_3	[10^{-3} , 10^{-2} , 10^{-1} , 10^0 , 10^1 , 10^2]
k	[5, 10 , 15, 20, 25, 30]

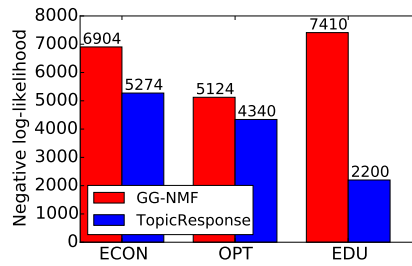


Fig. 5: Negative log-likelihood as goodness of fit; Smaller is better.

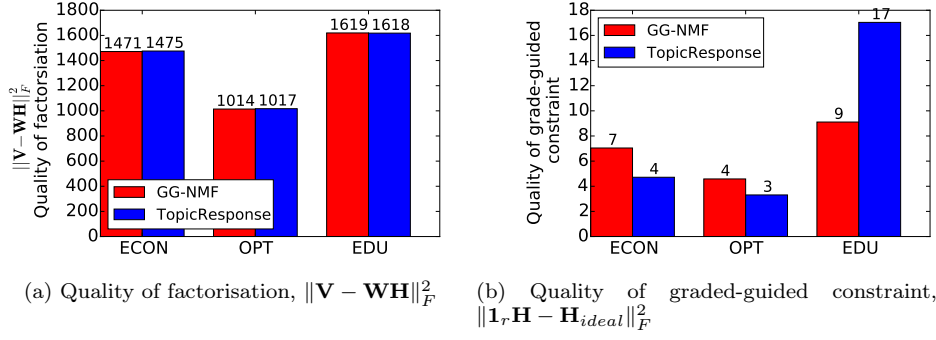


Fig. 6: Performance of GG-NMF and TOPICRESPONSE in terms of $\|\mathbf{V} - \mathbf{WH}\|_F^2$ and $\|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2$; Smaller is better.

soft constraint $\|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2$. For GG-NMF, the factorisation and Rasch estimation are separated, where topic-student response matrix \mathbf{H} is first obtained using GG-NMF, and then is taken as input to Rasch estimation. For TOPICRESPONSE, the negative log-likelihood is optimised together with factorisation. The parameters are set using the boldface default values in Table 4. Figure 5 displays the negative log-likelihood of GG-NMF and TOPICRESPONSE.

It can be seen that TOPICRESPONSE can yield superior negative log-likelihood, implying better fit between the topic-student response matrix \mathbf{H} and the Rasch model. TOPICRESPONSE therefore provides greater confidence that other item-level fit statistics such as infit, will be acceptable. *Jointly optimising the matrix factorisation and Rasch estimation can bring us closer to global optima.*

We present the results on quality of approximation $\|\mathbf{V} - \mathbf{WH}\|_F^2$ and supervision term $\|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2$, in Figure 6. From these plots, we can see that without sacrificing approximation performance in terms of $\|\mathbf{V} - \mathbf{WH}\|_F^2$, TOPICRESPONSE obtains superior $\|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2$ (while obtaining excellent negative log-likelihoods as above). This performance again demonstrates that optimising the factorisation and Rasch estimation globally can be superior to optimising them separately. We therefore conclude that TOPICRESPONSE is preferable to GG-NMF; we focus on results for TOPICRESPONSE in the remainder of our experiments.

5.5 Item Infit, Item Difficulty and Student Ability

We further examine the infit of each item, which indicates if the set of topics conform to the Rasch model, and is appropriate for measurement. As illustrated in Section 3.1.2, a conventional acceptable range of infit is 0.7 to 1.3. As an example, we show the item infit in Figure 7 on OPT MOOC. We can see that the infit of each item is in the acceptable range, with most very close to the (ideal) expected value of 1.0, indicating that the set of topics conform to the Rasch model and is appropriate for measuring student ability.

Additionally, we examine item difficulties and student abilities. Figure 8 displays the histogram of item difficulty and student ability along a common scale. According to the Rasch model, the higher a person's ability relative to the dif-

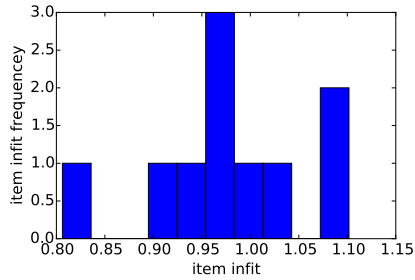


Fig. 7: Item infit histogram for OPT MOOC; infit closer to 1 is better.

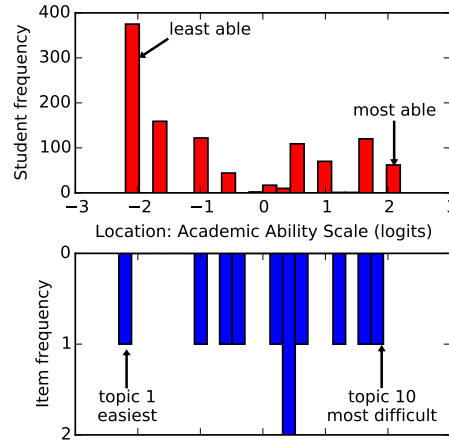


Fig. 8: Histograms of OPT MOOC student ability location (top) and item difficulty location (bottom).

faculty of a topic, the higher the probability that person posts on that topic. It can be seen that most students with low ability (around -2 logits), only dominate the “easiest” topic (topic 1 with difficult -2.3 logits); this topic concerns general problem solving. In other words, these students are likely to post only on topic 1, and unlikely to post on other topics. By comparison, the most able students with abilities around 2, with high probability contribute to all the topics.

5.6 Topic Interpretation and Discussion

We qualitatively examine topic interpretation, in order to assess educational meaningfulness. Well-scaled topics can potentially be used for curriculum refinement. Table 5 presents the topics generated using TOPICRESPONSE, alongside inferred difficulties. Topics are interpreted by an instructor who teaches a similar course. As the topics are not all course content-related, we envision that instructors examine discovered topics prior to using all for refining curriculum or taking other actions. Additionally, the inferred student ability levels and topic difficulty levels could be potentially used for personalised feedback, by tailoring appropriate topics of course content or forum discussion to students with their individual ability level taken into account. For example, most students (lowest ability) only discuss solving problem in general, as shown in Figure 8. If they cannot obtain sufficient help from forum discussions, they may be prone to drop out without further topic exploration. Therefore, in intervening with at-risk students, it is advisable to leverage discovered topics to better focus measures. Such services may be useful in preventing dropout in early stages (when most dropouts typically occur).

Table 5: Topics and difficulty levels, by TOPICRESPONSE on OPT MOOC.

No.	Topics	Interpretation	Inferred difficulty
1	use time problem get solut one optim algorithm tri work	Solving in general	-2.30
2	cours thank would lectur realli great assign good like think	Course feedback	-0.93
3	python use run program solver java matlab instal command work	Python/Java/Matlab (How to start)	-0.63
4	problem thank solut get grade knap- sack got feedback optim solv	How knapsack problem is solved and graded	-0.44
5	memori dp use column bb implement solv algorithm bound tabl	Comparing al- gorithms mem- ory/time	0.23
6	color node graph random edg greed opt search swap iter	Graph coloring	0.31
7	item valu weight capac estim take solut calcul best list	Knapsack problem	0.33
8	file pi line solver data submit lib urllib2 solveit open	Using solvers	0.52
9	video http class load lecture org prob- lem coursera optimization 001	Platform	1.17
10	submit assign assignment error messag view assignment_id detail class cours- era	Assignment sub- mission	1.73

5.7 Parameter Sensitivity

To validate the robustness of TOPICRESPONSE to parameter settings, a series of sensitivity experiments were conducted. The parameter settings are shown in Table 4. Negative log-likelihoods, $\|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2$, $\|\mathbf{V} - \mathbf{W}\mathbf{H}\|_F^2$ and $\|\mathbf{H} \circ \mathbf{H} - \mathbf{H}\|_F^2$ are examined in these experiments. Due to space limitations, we report here results for λ_0 on the OPT MOOC. The reader is referred to Appendix B for results on parameters $\lambda_1, \lambda_2, \lambda_3$, k on all three MOOCs.

Effect of Parameter λ_0 . As can be seen in Figure 9, as λ_0 is increased TOPICRESPONSE performs better in terms of negative log-likelihood, and performs worse in terms of the other three metrics due to the regularisation on the Rasch model. By contrast, the performance of GG-NMF does not change as there is no regularisation term on its Rasch estimation. Overall, TOPICRESPONSE performs well when λ_0 varies between 0.1 and 0.2.

6 Conclusion and Future Work

We have examined the suitability of content-based items (topics) discovered from MOOC forum discussions, for modelling student abilities. Our central tenet is that topics can be regarded as useful items for measuring latent skills, if student responses to these topics fit the Rasch item-response theory model, and if the discovered topics are further interpretable to domain experts. We propose to jointly optimise NMF and Rasch modelling, in order to discover Rasch-scaled topics. We

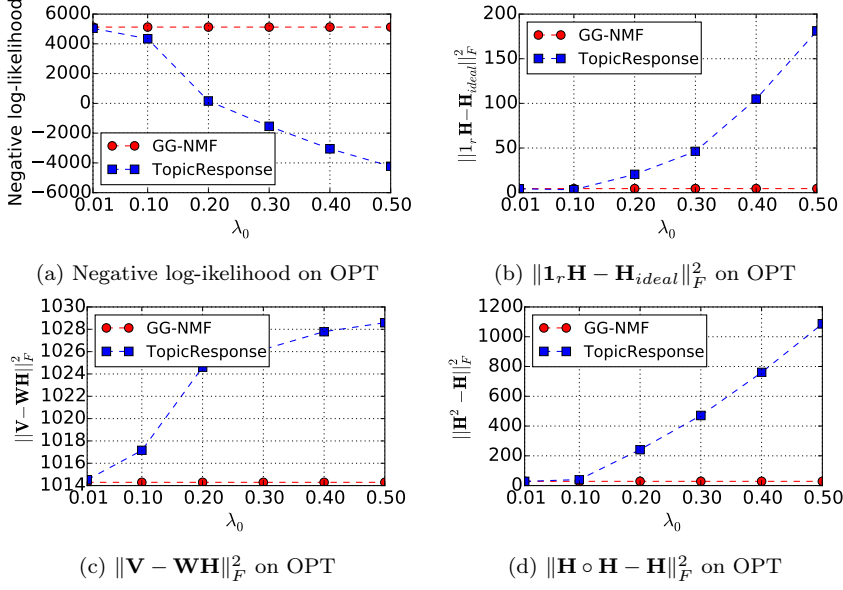


Fig. 9: Performance of GG-NMF and TOPICRESPONSE on OPT with varying λ_0 .

provide a quantitative validation on three Coursera MOOCs, demonstrating that TOPICRESPONSE yields better global fit to the Rasch model (observed with lower negative log-likelihood), maintains good quality of factorisation approximation, while measuring the students' academic abilities (reflected by the grade-guided constraint on students' participation on topics). We also provide qualitative examination of topic interpretation with inferred difficulty levels on a Discrete Optimisation MOOC. The results on goodness of fit and our qualitative examination, together suggest potential applications in curriculum refinement, student assessment and personalised feedback.

We opted to study the relatively simple Rasch model, as it forms the basis of very many subsequent models in the literature. One direction for extension, is that for any model (like Rasch), that fits parameters via maximum-likelihood estimation (or risk minimisation in general), the model can be augmented with NMF as an additional regularisation. For example, such an extension should be straightforward for polychotomous observations, hierarchical models on latent skills, models that include more flexible per-student variation, etc. These represent fruitful direction for future research. Another possible extension could involve augmenting the \mathbf{W}, \mathbf{H} matrices in the NMF or Rasch objective terms with manually-crafted items, to make effective use of prior knowledge.

Acknowledgements We thank Jeffrey Chan for discussions related to this work, and the anonymous reviewers and editor for their thoughtful feedback. This work is supported by Data61, and the Australian Research Council (DE160100584).

References

- Bachrach Y, Graepel T, Minka T, Guiver J (2012) How to grade a test without knowing the answers—a Bayesian graphical model for adaptive crowdsourcing and aptitude testing. In: Proceedings of the 29th International Conference on Machine Learning (ICML-12), pp 1183–1190
- Baker FB, Kim SH (2004) Item response theory: Parameter estimation techniques. CRC Press
- Bergner Y, Droschler S, Kortemeyer G, Rayyan S, Seaton D, Pritchard DE (2012) Model-based collaborative filtering analysis of student response data: Machine-learning item response theory. International Educational Data Mining Society
- Bertsekas DP (1999) Nonlinear programming. Athena Scientific
- Bond TG, Fox CM (2001) Applying the Rasch model: Fundamental measurement in the human sciences. Lawrence Erlbaum Associates Publishers
- Champaign J, Colvin KF, Liu A, Fredericks C, Seaton D, Pritchard DE (2014) Correlating skill and improvement in 2 MOOCs with a student’s time on tasks. In: Proceedings of the First ACM Conference on Learning@Scale Conference, ACM, pp 11–20
- Chaturvedi S, Goldwasser D, Daumé III H (2014) Predicting instructor’s intervention in MOOC forums. In: ACL (1), pp 1501–1511
- Chen CM, Lee HM, Chen YH (2005) Personalized e-learning system using item response theory. Computers & Education 44(3):237–255
- Colvin KF, Champaign J, Liu A, Fredericks C, Pritchard DE (2014) Comparing learning in a MOOC and a blended on-campus course. In: Educational Data Mining 2014
- Gillani N, Eynon R, Osborne M, Hjorth I, Roberts S (2014) Communication communities in MOOCs. arXiv preprint arXiv:14034640
- Guttman L (1950) The basis for scalogram analysis. In: Stouffer S (ed) Measurement and Prediction: The American Soldier, Wiley, New York
- He J, Rubinstein BI, Bailey J, Zhang R, Milligan S, Chan J (2016) MOOCs meet measurement theory: A topic-modelling approach. In: Thirtieth AAAI Conference on Artificial Intelligence
- Jenders M, Krestel R, Naumann F (2016) Which answer is best? Predicting accepted answers in MOOC forums. WWW’2016 Companion
- Lee DD, Seung HS (1999) Learning the parts of objects by non-negative matrix factorization. Nature 401(6755):788–791
- Lee DD, Seung HS (2001) Algorithms for non-negative matrix factorization. In: Advances in Neural Information Processing Systems, pp 556–562
- Linacre JM (2002) What do infit and outfit, mean-square and standardized mean. Rasch Measurement Transactions 16(2):878
- Linacre JM (2006) Misfit diagnosis: Infit outfit mean-square standardized. Retrieved June 1:2006
- Milligan S (2015) Crowd-sourced learning in MOOCs: Learning analytics meets measurement theory. In: Proceedings of the Fifth International Conference on Learning Analytics And Knowledge, ACM, pp 151–155
- Ramesh A, Kumar SH, Foulds J, Getoor L (2015) Weakly supervised models of aspect-sentiment for online course discussion forums. In: Annual Meeting of the Association for Computational Linguistics (ACL)

- Rasch G (1993) Probabilistic models for some intelligence and attainment tests. ERIC
- Scholten AZ (2011) Admissible statistics from a latent variable perspective. *Theory & Psychology* 18:111–117
- Wen M, Yang D, Rose C (2014) Sentiment analysis in MOOC discussion forums: What does it tell us? In: *Educational Data Mining 2014*
- Wright BD, Masters GN (1982) Rating Scale Analysis. Rasch Measurement. ERIC
- Wright BD, Linacre JM, Gustafson J, Martin-Lof P (1994) Reasonable mean-square fit values. *Rasch measurement transactions* 8(3):370
- Yang D, Adamson D, Rosé CP (2014) Question recommendation with constraints for massive open online courses. In: *Proceedings of the 8th ACM Conference on Recommender systems*, ACM, pp 49–56
- Yang D, Wen M, Howley I, Kraut R, Rose C (2015) Exploring the effect of confusion in discussion forums of massive open online courses. In: *Proceedings of the Second (2015) ACM Conference on Learning@ Scale*, ACM, pp 121–130
- Zhang Z, Ding C, Li T, Zhang X (2007) Binary matrix factorization with applications. In: *Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on*, IEEE, pp 391–400
- Zhang ZY, Li T, Ding C, Ren XW, Zhang XS (2010) Binary matrix factorization for analyzing gene expression data. *Data Mining and Knowledge Discovery* 20(1):28–52

A Proof of Theorem 1

The update rules for β_i and θ_j are derived using the Newton-Raphson method (Bertsekas, 1999), where the convergence to a local optimum is guaranteed. Here, we focus on the proof for the update rule for h_{ij} . The update rule for w_{ij} can be proved similarly. We closely follow the procedure described in (Lee and Seung, 2001), where an auxiliary function similar to that used in the Expectation-Maximization (EM) algorithm is used for proof.

Definition 2 (Lee and Seung 2001) $G(h, h')$ is an auxiliary function for $F(h)$ if the conditions

$$G(h, h') \geq F(h) \quad , \quad G(h, h) = F(h) \quad ,$$

are satisfied.

Lemma 1 (Lee and Seung 2001) If G is an auxiliary function, then F is non-increasing under the update

$$h_{t+1} = \underset{h}{\operatorname{argmin}} G(h, h^t) \quad . \quad (18)$$

Proof The result follows from noting $F(h^{t+1}) \leq G(h^{t+1}, h^t) \leq G(h^t, h^t) = F(h^t)$. \square

For any element h_{ij} in \mathbf{H} , let $F_{h_{ij}}$ denote the part of $f(\mathbf{W}, \mathbf{H}, \boldsymbol{\theta}, \boldsymbol{\beta})$ in Eq. (12) relevant to h_{ij} . Since the update is essentially element-wise, it is sufficient to show that each $F_{h_{ij}}$ is non-increasing under the update rule of Eq. (15). To prove this, we define the auxiliary function regarding h_{ij} as follows.

Lemma 2 *Function*

$$G(h_{ij}, h_{ij}^t) = F_{h_{ij}}(h_{ij}^t) + F'_{h_{ij}}(h_{ij}^t)(h_{ij} - h_{ij}^t) + \varphi_{ij}(h_{ij} - h_{ij}^t)^2 \quad , \quad (19)$$

where

$$\varphi_{ij} = \frac{2(\mathbf{W}^T \mathbf{W} \mathbf{H})_{ij} + 2\lambda_1(\mathbf{1}_r^T \mathbf{1}_r \mathbf{H})_{ij} + 12\lambda_2(h_{ij}^t)^3 + 2\lambda_2 h_{ij}^t + \lambda_3(\boldsymbol{\theta} - \boldsymbol{\beta})_{ij}^-}{2h_{ij}^t}$$

is an auxiliary function for $F_{h_{ij}}$.

Proof It is obvious that $G(h_{ij}, h_{ij}) = F_{h_{ij}}$. So we need only prove that $G(h_{ij}, h_{ij}^t) \geq F_{h_{ij}}$. Considering the Taylor series expansion of $F_{h_{ij}}$,

$$F_{h_{ij}} = F_{h_{ij}}(h_{ij}^t) + F'_{h_{ij}}(h_{ij}^t)(h_{ij} - h_{ij}^t) + \frac{1}{2}F''_{h_{ij}}(h_{ij}^t)(h_{ij} - h_{ij}^t)^2,$$

$G(h_{ij}, h_{ij}^t) \geq F_{h_{ij}}$ is equivalent to $\varphi_{ij} \geq \frac{1}{2}F''_{h_{ij}}(h_{ij}^t)$, where

$$F''_{h_{ij}}(h_{ij}^t) = 2(\mathbf{W}^T \mathbf{W})_{ii} + 2\lambda_1(\mathbf{1}_r^T \mathbf{1}_r)_{ii} + 12\lambda_2(h_{ij}^t)^2 - 12\lambda_2 h_{ij}^t + 2\lambda_2.$$

To prove the above inequality, we have

$$\begin{aligned} \varphi_{ij} h_{ij}^t &= (\mathbf{W}^T \mathbf{W} \mathbf{H})_{ij} + \lambda_1(\mathbf{1}_r^T \mathbf{1}_r \mathbf{H})_{ij} + 6\lambda_2(h_{ij}^t)^3 + \lambda_2 h_{ij}^t + 0.5\lambda_3(\boldsymbol{\theta} - \boldsymbol{\beta})_{ij}^- \\ &= \sum_{l=1}^k (\mathbf{W}^T \mathbf{W})_{il} h_{lj}^t + \lambda_1 \sum_{l=1}^k (\mathbf{1}_r^T \mathbf{1}_r)_{il} h_{lj}^t + 6\lambda_2(h_{ij}^t)^3 + \lambda_2 h_{ij}^t + 0.5\lambda_3(\boldsymbol{\theta} - \boldsymbol{\beta})_{ij}^- \\ &\geq (\mathbf{W}^T \mathbf{W})_{ii} h_{ij}^t + \lambda_1(\mathbf{1}_r^T \mathbf{1}_r)_{ii} h_{ij}^t + 6\lambda_2(h_{ij}^t)^3 + \lambda_2 h_{ij}^t - 12\lambda_2 h_{ij}^t \\ &\geq h_{ij}^t ((\mathbf{W}^T \mathbf{W})_{ii} + \lambda_1(\mathbf{1}_r^T \mathbf{1}_r)_{ii} + 6\lambda_2(h_{ij}^t)^2 - 6\lambda_2 h_{ij}^t + \lambda_2) \\ &= \frac{1}{2}F''_{h_{ij}}(h_{ij}^t) h_{ij}^t. \end{aligned}$$

Thus, $G(h_{ij}, h_{ij}^t) \geq F_{h_{ij}}$ as claimed. \square

Replacing $G(h_{ij}, h_{ij}^t)$ in Eq. (18) by Eq. (19) and setting $\frac{\partial G(h_{ij}, h_{ij}^t)}{\partial h_{ij}}$ to be 0 results in the update rule in Eq. (15). Since Eq. (19) is an auxiliary function, $F_{h_{ij}}$ is non-increasing under this update rule.

B Experimental Results of Parameter Sensitivity on Regularisation Parameters $\lambda_1, \lambda_2, \lambda_3$, and k

a) Effect of Parameter λ_1 : As we can see from Figure 11, GG-NMF and TopicResponse are not sensitive to λ_1 , performing stably with varying λ_1 . TopicResponse constantly performs better in terms of negative log-likelihood while maintaining the comparable performance in terms of the other three metrics.

b) Effect of Parameter λ_2 : It can be seen from Figure 12 that GG-NMF and TopicResponse perform well in terms of $\|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2$ (Figure 12d to Figure 12f) and $\|\mathbf{H} \circ \mathbf{H} - \mathbf{H}\|_F^2$ (Figure 12j to Figure 12l) when λ_2 varies from 10^0 to 10^2 , and from 10^{-3} to 10^0 respectively. $\|\mathbf{V} - \mathbf{W} \mathbf{H}\|_F^2$ gets worse as λ_1 increases, but does not change a lot compared to $\|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2$ and $\|\mathbf{H} \circ \mathbf{H} - \mathbf{H}\|_F^2$. As λ_2 increases, the performance of GG-NMF and TopicResponse in terms of negative log-likelihood decrease, and TopicResponse constantly performs better than GG-NMF. Overall, λ_2 with values around 1.0 is good for GG-NMF and TopicResponse.

c) Effect of Parameter λ_3 : It can be seen that GG-NMF and TopicResponse perform well in terms of $\|\mathbf{1}_r \mathbf{H} - \mathbf{H}_{ideal}\|_F^2$ (Figure 13d to Figure 13f) and $\|\mathbf{H} \circ \mathbf{H} - \mathbf{H}\|_F^2$ (Figure 13j to Figure 13l) when λ_3 varies from 10^{-1} to 10^0 , and from 10^0 to 10^2 respectively. Similar to λ_2 , λ_3 does not affect $\|\mathbf{V} - \mathbf{W} \mathbf{H}\|_F^2$ significantly. TopicResponse constantly achieves better negative log-likelihood than GG-NMF. Overall, λ_3 with values around 1.0 is good for GG-NMF and TopicResponse.

d) Effect of the number of topics k : It can be seen from Figure 14 that TopicResponse constantly outperforms GG-NMF in terms of negative log-likelihood, while getting slightly worse performance in the other three metrics. This is reasonable, as GG-NMF has more constraints and hence the model itself is less likely to perform as well as the less constrained GG-NMF in other metrics. Overall, GG-NMF and TopicResponse perform well in the experiments when k is set to 10 or 15. We choose 10 as the value of k since a smaller number of topics are easier to analyse.

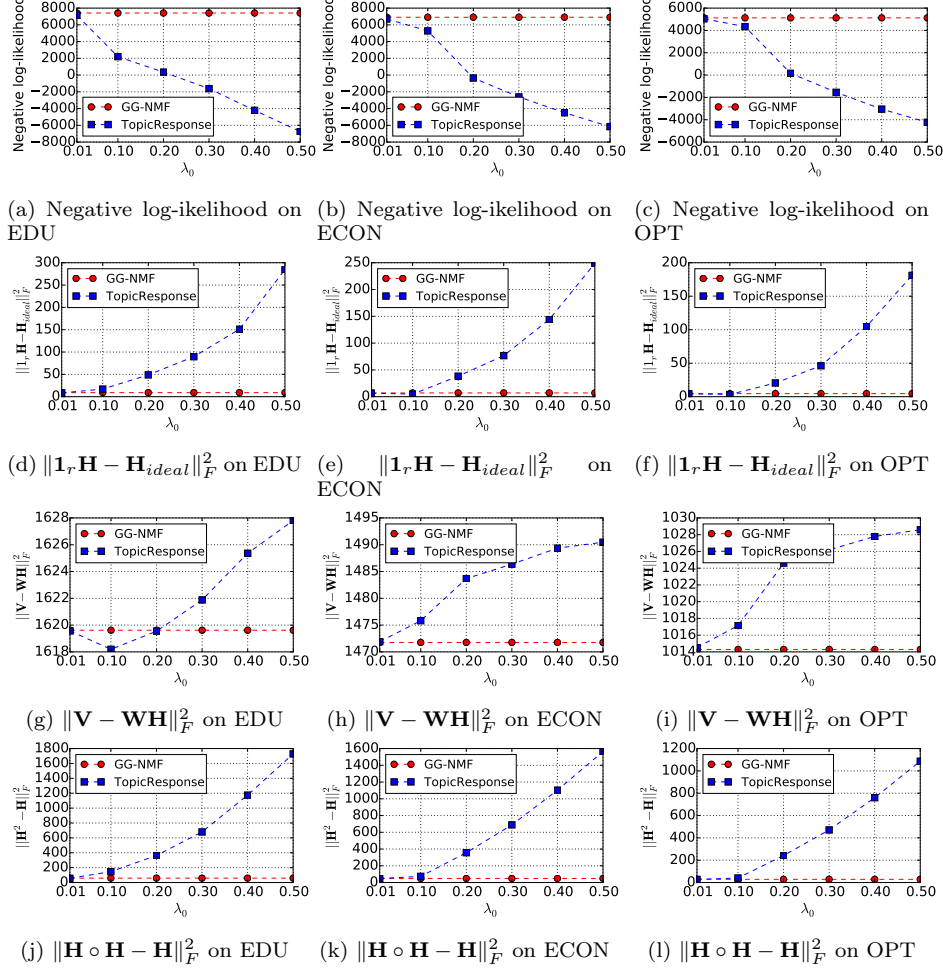
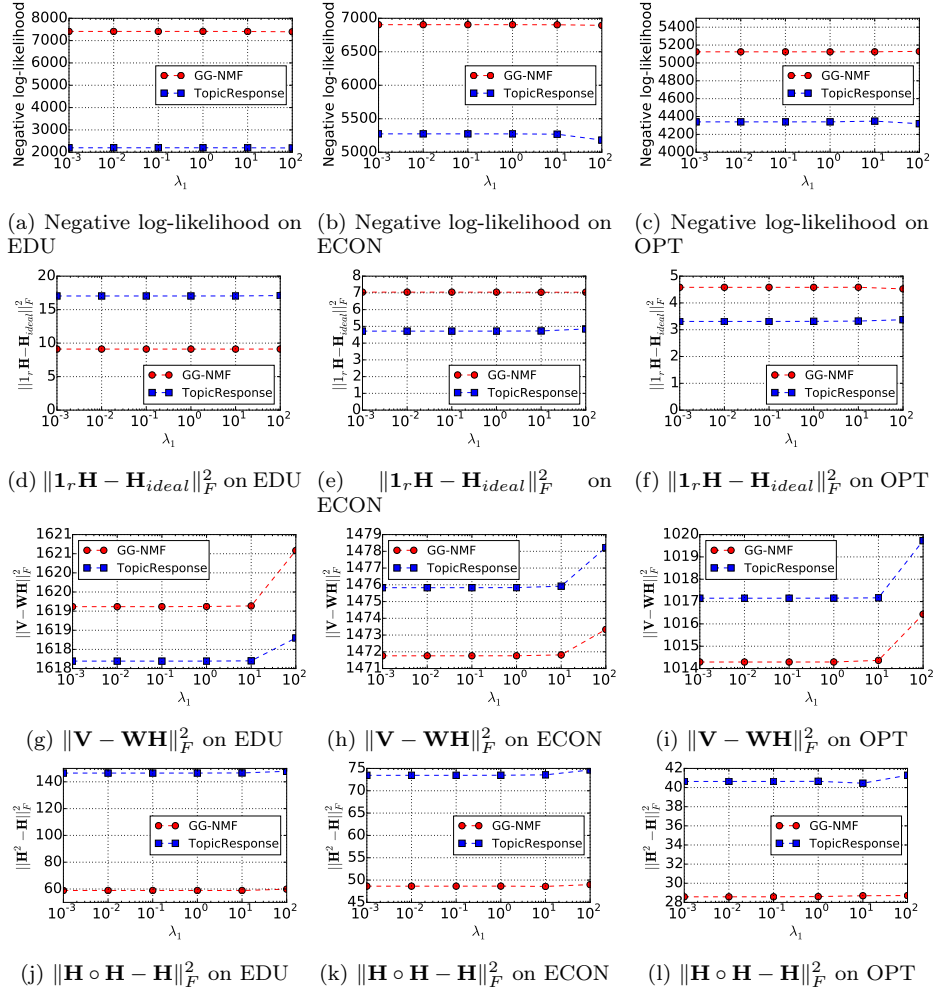
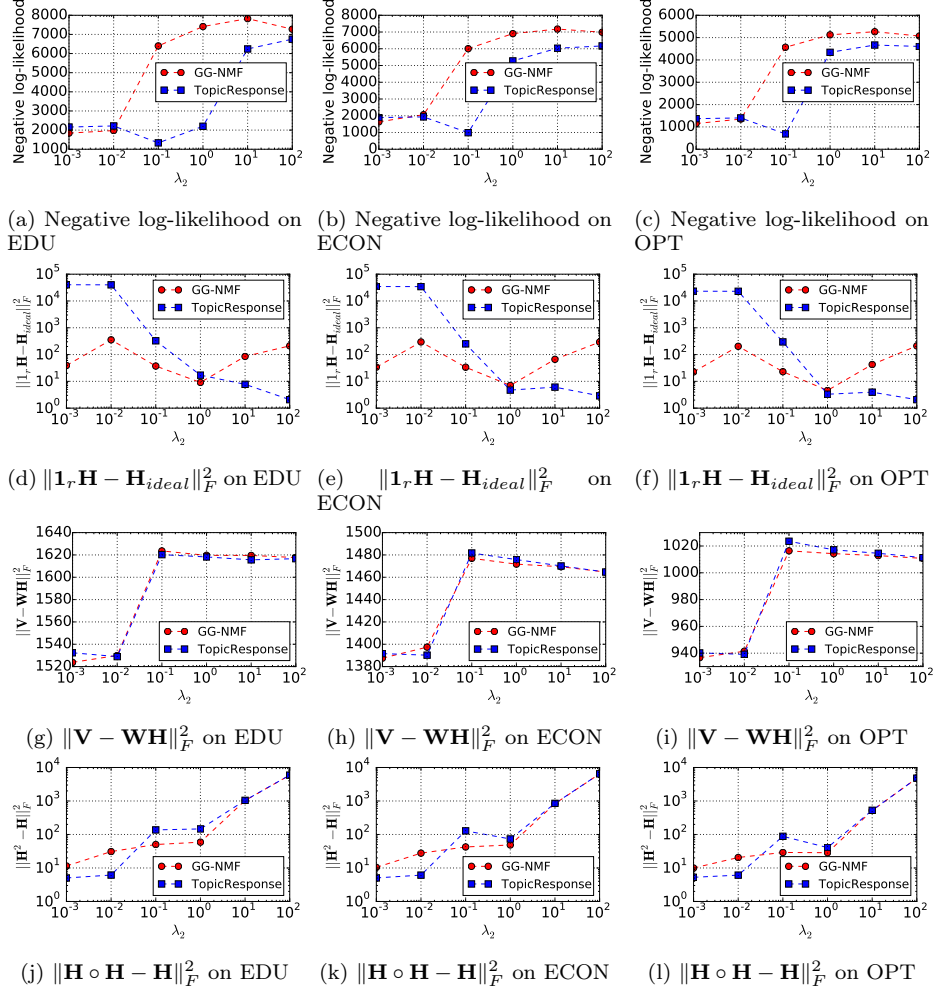
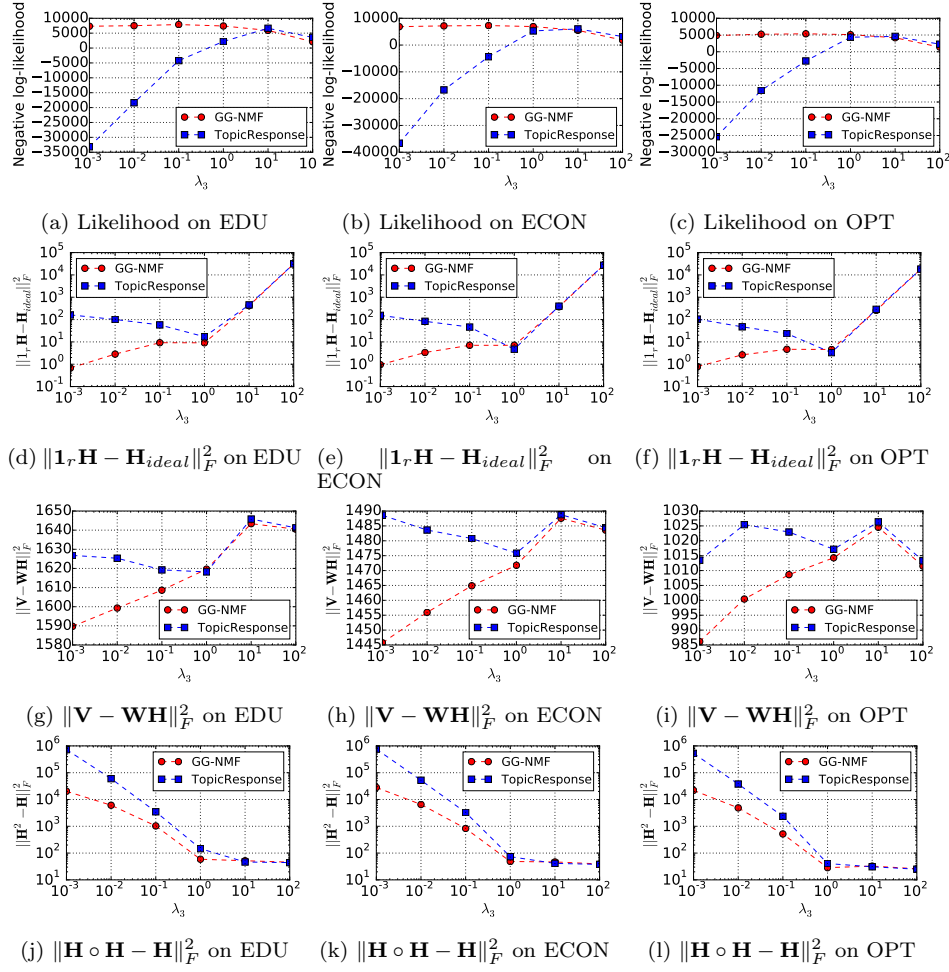
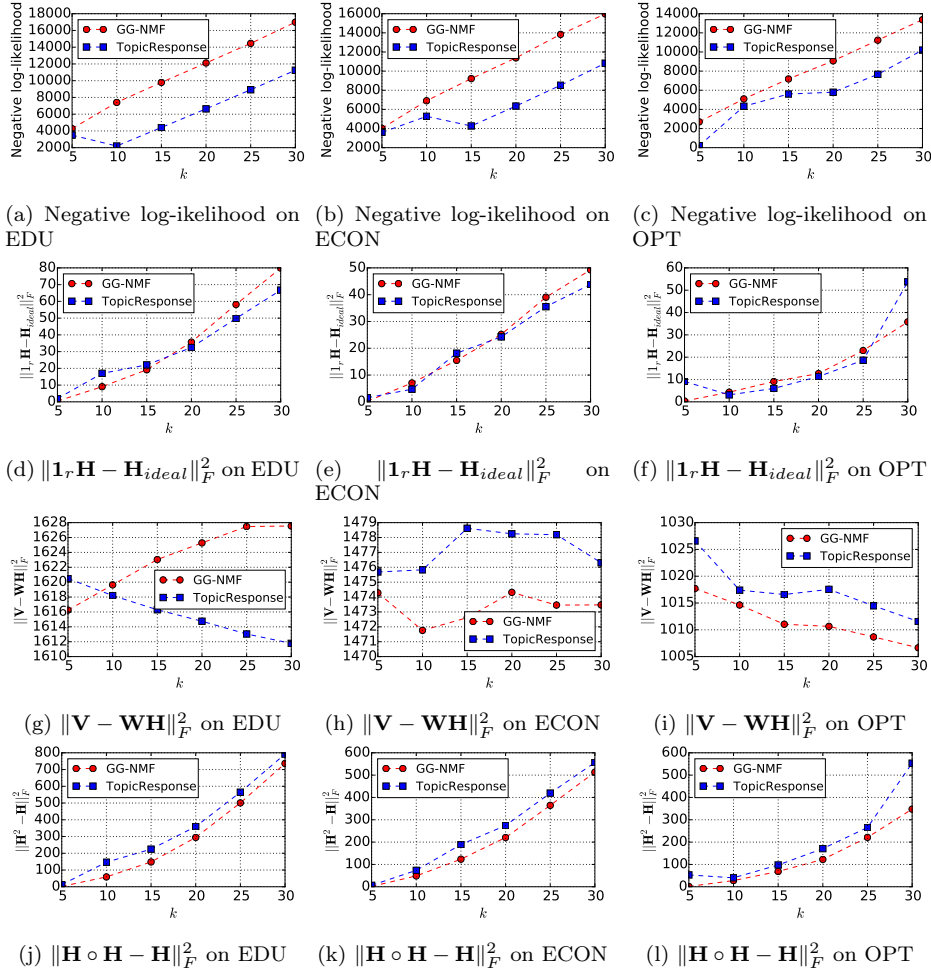


Fig. 10: Performance of GG-NMF and TopicResponse with varying λ_0 .

Fig. 11: Performance of GG-NMF and TopicResponse with varying λ_1 .

Fig. 12: Performance of GG-NMF and TopicResponse with varying λ_2 .

Fig. 13: Performance of GG-NMF and TopicResponse with varying λ_3 .

Fig. 14: Performance of GG-NMF and TopicResponse with varying k .